The Game of Prediction and Retrodiction: A Radical-Constructivist Perspective on the Notion of Time in Physics

ANDREAS QUALE
University of Oslo

ABSTRACT: The notions of *prediction* and *retrodiction*, and the role they play in the natural sciences, are discussed. These notions derive from our perception of the fundamental category of *time*, as an ordering scheme for our experiential world. The issue of *philosophical determinism vs. human free will* is examined from a perspective of radical constructivism, and contrasted with the issue of *solipsism vs. shared experience*; and it is argued that both philosophical determinism and solipsism may be rejected, on the same (existential, not logical) grounds. Both prediction and retrodiction are discussed, in the context of some sciences (notably, classical and quantum physics), and are shown to be realisable only to a very limited extent. Some consequences of this, for the ability of science to forecast what will happen in the future, or to infer what has happened in the past, are reviewed. It is concluded that scientific knowledge, of both the past and the future, is (and must be) *constructed in the present*, based on presently observed data and theoretical arguments. Hence there can be no conception of *true knowledge*, either of the future or of the past.

RESUME: Les notions des mots *prédiction* et *rétrodiction* ainsi que le rôle qu'ils jouent dans les sciences naturelles viennent ici animer le débat. Ces notions viennent de la perception que nous avons du concept vital du *temps;* mécanisme qui commande notre monde concret. La question sur le *déterminisme philosophique par opposition à la liberté des hommes* sera examinée à partir d'une perspective sur le constructivisme radical et contrastée ainsi que la question sur le *solipsisme par opposition à l'expérience partagée.* De surcroît, il est question que le déterminisme et le solipsisme philosophiques peuvent être, tous deux, rejetés pour les mêmes motifs (existentiels et non logiques). Toutes deux, la prédiction et la rétro diction font l'objet du débat dans le contexte de certaines sciences (notamment dans le domaine de la physique classique et quantique) et sont présentées comme étant réalisables uniquement

Journal of Educational Thought
Vol. 42, No. 2, 2008, 105-125.

> dans un champ très limité. Certaines conséquences tirées de l'habilité des sciences à prévoir l'avenir ou à déduire ce qui est arrivé dans le passé, sont reconsidérées. En conclusion, la connaissance scientifique des deux, passé et avenir, est (et doit être) *construite dans le présent*, fondée sur des données observées dans l'instant et sur des discussions théoriques. Ainsi donc, le concept de la *connaissance réelle* n'existe peut-être pas, que ce soit celle de l'avenir ou celle du passé.

Some aspects of the category of time are addressed, as they appear in the formulation of fundamental natural science – in particular, of physics. The perspective is that of radical constructivism, as formulated by Ernst von Glasersfeld (1981, 1984, 1989, 1992, 1995, 2000), which is based on two propositions:

- Knowledge is learnt through active construction by the knower, not by passive reception and storage of data in the mind;
- The purpose of knowledge construction is to organise the experiential world of the knower – it does not (and cannot) reflect any objective truth about the world.

In the mind of the public, one of the most remarkable achievements of modern natural science is its ability to *foresee the future*: that is, to predict events that are going to happen, before they actually do happen. Announcements in the media that there will be an eclipse of the Sun (or the Moon), or that a comet will become visible in the sky, always excite great interest; and when it then turns out that these phenomena occur precisely as predicted, this never fails to impress. Indeed, there is evidence that many people think of scientific prediction as a kind of modern magic – except that it is generally assumed to be much more accurate and reliable than the methods of the old magicians and fortune-tellers, with their incantations, crystal balls, and tarot cards. Moreover, it is universally acknowledged that this predictive ability of science can be put to actual use, designing and building technological devices that will indeed work in the way that scientific theory tells us they will.

It should be born in mind, of course, that this gift of prophecy is not always to be trusted. For example, everybody knows that weather forecasting is not reliable: in fact, it may well rain tomorrow, even though the forecast is for clear skies! At least two popular explanations for this failure may then be encountered: (a) predicting the weather is a difficult task, which the scientists have not yet mastered completely (though they surely will, given time); or (b) the phenomenon of weather exhibits an inherent capriciousness, and hence cannot be predicted accurately by scientific means – there will always be some element of

guesswork involved. In any case, this does not shake the general conviction that science is, to a large extent, about *controlling the future*: either through passive prediction of coming events, or through active manipulation of technology, as indicated above. Indeed, one might say that the whole conception of science as a *description of observed regularities in the world* leads quite naturally to the notion of prediction: that is, to an idea of what will happen in the future, inspired by the experience of a regular pattern of events as observed up to now.

On closer inspection, however, things turn out to be a little more complicated. First, it should be noted that many sciences are less concerned with *prediction* than with *retrodiction* – in other words, they address issues of "what did happen in the past," rather that "what will happen in the future." In science education, the prime example is of course evolutionary biology: here, conditions in the past and the development of species in early times are inferred from (fossil and other) evidence that is observed today.

Moreover, some sciences can in certain particular cases do both: that is, use observations made in the present to predict the future *and* retrodict the past. Thus, for instance, Newtonian mechanics allows us to calculate, not only when the next solar eclipse will occur, but also when the last one did occur. This, however, is a special case: there are in fact many physical systems that do not exhibit such symmetry between past and future – that is, a capability to both predict and retrodict events. We will return to this point later.

Before going on, it is important to note that the notions of prediction and retrodiction, as used here, are both defined relative to a *theoretical framework*; in other words, they require not only an act of observation, but also the use of a theory. One simple example: based on day-to-day experience, any individual person can prophesy with confidence that the sun will rise tomorrow morning and in subsequent future mornings – and similarly assume that it has done so regularly in past mornings, since long before she or he was born. But this kind of prophecy or assumption (based on inference from past experience) is not deduced from a theory, that offers an explanation of *why* such a regular behaviour is in fact observed; and hence it does not qualify as prediction (or retrodiction), in the present context.

In this paper, I discuss some issues connected with prediction and retrodiction in science, in the context of radical constructivism. But first, we need to examine closer the notion of time.

A Perspective on Time

Any construction of cognitive knowledge about our experiential world must involve the notions of *time* and *space*. We cannot conceive of anything as happening, unless it is spatially and temporally situated – that is, unless it occupies some location in space and some period of time.

Before the 20th century it was customary to think of these two items (time and space) as just being there, forming a constant background – an empty stage, as it were, on which the action of the world is played out. With the advent of modern physics, however, this view has become more problematic:

- In special relativity theory, the measure of space and time depend on the motion of the observer;
- In general relativity theory, space and time enter into the dynamical equations, as physical fields (describing the space-time geometry) that interact with other physical entities;
- The development of quantum physics has further eroded our confidence in a fixed background of space and time, which is unaffected by events taking place "in the foreground."

In the perspective of cognitive learning, it is more natural to consider time and space as *experiential categories*, leaving aside any speculations as to whether they can meaningfully be ascribed a physical existence. This is the viewpoint adopted by radical constructivism: *we* (the learners) use our cognitive faculties to situate our experiences in time and space, in the construction of our knowledge about them – as was, in fact, also posited by Immanuel Kant, in the 18th century.

Thus, for instance, we may construct the notion of *object permanence* under displacements in time and space: *the chair* I see by the window in front of me now is *the same* chair as the one I saw there yesterday; and when I move it to another room, it will still be *the same* chair. (Note, however, that this notion of permanence is not wholly unproblematic in modern physics: for instance, quantum theory holds that a system does not exist in a definite state, with observable physical properties, before it is actually observed. We will get back to that.)

In the present context of prediction and retrodiction, the important experiential category is that of *time*. Now, the nature of time, and the role it plays in human cognition and natural science, is a topic that has fascinated philosophers for a long time. It provides a fundamental basis for our entire perception of the world, and is perhaps for that reason rather difficult to come to grips with; indeed, the 4th-century philosopher St. Augustine is reported to have said something to the effect that "if you

do not ask me what time is, I know it; but if you ask me, then I do not know."

So, let us examine the category of time, as it may be experienced by an individual learner. This perspective may be illustrated:

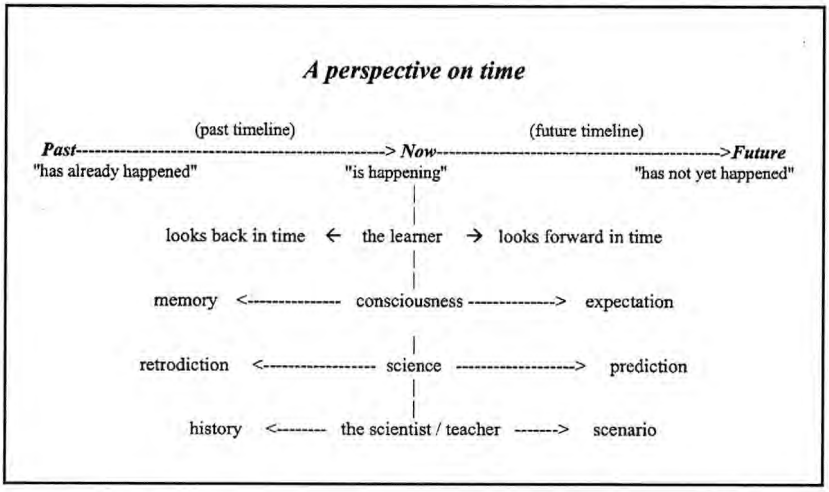

Every individual learner can conceive of him or herself as situated in a moment of time *Now*, on a timeline extending into their own *Past* and *Future*. On the past timeline are located all events that have already happened before *Now*, and on the future timeline are located all events that have not yet happened up to *Now*.

This time-line is thought of as being *directed* from past to future, in the sense that events on it are ordered in a before-after relationship, represented in the figure as a *left-right* positioning of points on the line. Thus, for any two events *A* and *B*, located in points on the timeline, one of the following three must be the case: either *A* happens *before* (lies to the left of) *B*, or *A* happens *after* (lies to the right of) *B*, or *A* and *B* are *simultaneous* (occupy the same point on the line). It follows that all events on the past timeline are before *Now*, and all events on the future timeline are after *Now*.

(It may be remarked that this looks rather like the structuring of physical events, as points on a time-like observer world-line, in the four-dimensional space-time featured in the theories of special and general relativity. It should be emphasised, however, that there is no pretence of mathematical rigour here: the picture is intended to describe the learner's *intuitive experience* of her or his own world, as segmented into past, present, and future events.)

Let us now examine the notions of prediction and retrodiction, with reference to this perspective of time. We may conceive of these acts as performed by a learner, situated in the *Now* moment – in other words, as carried out by some person located in the present – who is then looking either *forward* in time (towards her/his future) or *backward* in time (towards his or her past).

For instance, we may consider the consciousness of the individual learner. From the vantage point of *Now*, his or her conscious mind then has the following two windows into time: her or his *expectations* (of what will happen in the future), and her or his *memories* (of what has happened in the past).

Alternatively, adopting a more formalised point of view: a science may feature a theory which attempts to *predict* what will happen in the future, and/or to *retrodict* what did happen in the past, both with reference to conditions that are observed in the present.

And the scientist or teacher, in presenting such a theory, will typically do so by outlining a *scenario* (of what conditions will be in the future), or a *history* (of what conditions were in the past). Concrete examples of this will be discussed presently.

There is an asymmetry between the past and the future, with respect to their degree of determinism. The past is generally taken to be fixed: events on the past timeline have already happened, and cannot be changed. But the events of the future have not yet happened, and the question arises: to what extent is it possible, for the individual at the moment *Now*, to affect the course of such events – to be master of her or his own future, so to speak? This, of course, touches on the deep issue of philosophical determinism: do we, as human beings, possess a *free will*, and thus the ability to control (or at least exert some influence on) our own future life? Or, are future events *totally determined causally* by past events, so that any freedom we may feel that we have, in choosing between different courses of action in the present, is illusory?

This latter viewpoint may, in fact, be regarded as the "reverse of solipsism." In a solipsist scenario, the world has no effect on our construction of it, so that we can dream it up to be whatever we want; whereas in a totally deterministic framework, the world is pre-set to behave in a certain manner, which we have no means of influencing!

The stance of philosophical determinism is rejected by radical constructivism, on the same grounds that solipsism is rejected. One can certainly argue that a philosophical-determinist position is *logically irrefutable*. Although it appears to us that we have the freedom to choose between alternatives, there is no way we can actually prove that this is

not just an illusion – that is, that our own choice of any particular course of action is not, in fact, totally predetermined by past events. (Similarly, one cannot logically disprove the solipsist proposition: that our whole experiential world is just a figment of our imagination, so that in fact nothing exists outside our individual minds!)

However, such a position is *existentially irrelevant* – it simply does not play any part in the way that we actually choose to conduct our lives! In fact, our experience is that we are continually faced with choices every waking moment of our lives; and we then feel that we are free to choose, and act accordingly. (Similarly, we disregard the solipsist stance in the interpretation of our individual experiences, acting instead on the tacit premise that these experiences do indeed refer to a shared experiential world.)

Thus, radical constructivism (like virtually every other epistemic position of relevance for education) accepts as a fundamental premise that we do have freedom to act, and that our construction of knowledge takes place in the interplay between our own free will and the unfolding of natural processes in time as we perceive them.

As we shall see, a theory describing these processes will sometimes allow for future events to be determined by events preceding them. But even then, human free will still rules *in the present*:

- a scientist is free to focus on whatever scientific field she or he may find to be of interest;
- an experimenter may select and set up the initial conditions of her or his experiment in many different ways, according to what kind of outcome that person is looking for;
- a designer may choose between different implementations of whatever device that person is designing, depending on the purpose, intentions, personal tastes, economic constraints, and so forth.

Thus, we shall have to distinguish between two kinds of determinism: *physical determinism*, where natural laws compel future conditions to be fixed by conditions in the past, while the human will is still free to act on conditions in the present; and the much stronger position (rejected here) of *philosophical determinism*, where everything that happens in the future (including human volition and action) is totally determined by events in the past.

Returning to the timeline figure, we note that the individual learner who is situated in the moment *Now* may conceptualise his or her relationship with the events on the timeline, in at least three different ways:

1. The learner may imagine that she or he is stationary, and that *time flows past her/him*, from right to left on the timeline. This is the origin of figures of speech like "the passage of time," "as time goes by," "coming events," and so forth. The image invoked is one in which events approach her or him from the future, happen to her or him *Now*, and then recede away into the past.

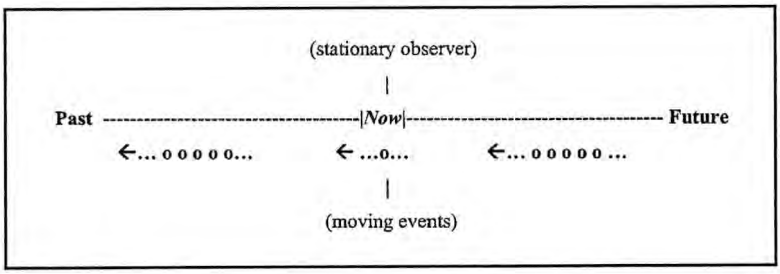

2. The learner may conceive of the events on the line as being at rest, and herself or himself (or, equivalently, the moment *Now*) as moving through, from left to right on the timeline. In other words, the individual may think of oneself as "leaving the past behind" and "heading towards the future."

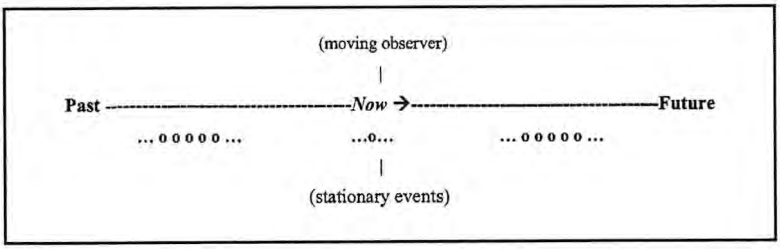

3. The learner may regard the events on the time-line as forming a path in a four-dimensional space-time: that is, a *timelike observer world-line*, in the terminology of relativity theory. In this space-time, there is no conception of "motion through time" – worldlines are static geometric objects.

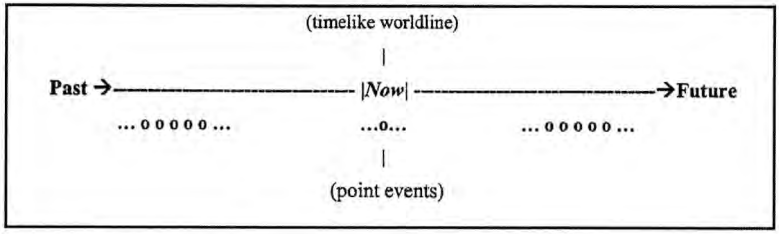

This leads to the notion of *time windows*, as considered in the perspective of prediction and retrodiction:

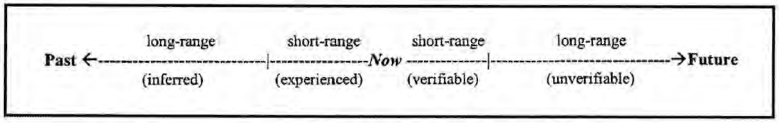

The *future time window* of the learner, as seen from the moment ***Now***, may be taken to consist of a short-range and a long-range segment. The short-range segment comprises those future events that are in principle open to verification by the learner acting *Now*: by performing experiments to test some prediction of a theory, or (more indirectly) by designing devices that work as predicted by the theory. The long-range segment contains events that will occur so far into the future that any prediction of them cannot be verified by this individual.

We will refer to these two segments as the *near future* and the *distant future*. There is, of course, no sharp divide envisaged between the two; but we may imagine the extent of the near future to be of the order of a human lifetime.

A similar segmentation may be performed for the *past time window*. We can imagine it to be divided into two segments of past events:

- a short-range segment, containing the *near past* of the learner *Now*. This consists of events that have been experienced by the individual in her or his own past, as accessed through personal memory. Here we shall also include concurrent events, which have been experienced by other individuals in the same period of time, and of which the learner has knowledge – in other words, events that she or he in principle could have experienced, if he or she had been there when they happened. Indeed, events in this segment are often referred to as having occurred "within living memory."

- a long-range segment, containing the *distant past* of the individual. These are events that the learner could not possibly have experienced herself, because they happened well before he or she was born. The characteristics of these events will then have to be inferred by that individual, using whatever evidence (hearsay, written documents, artefacts, fossils) that may be at hand.

Here again, there is no sharp divide between the two segments; but the extent of the immediate past may be thought of as being of the order of a human lifetime.

Prediction and Retrodiction in Physics

The prime example of a predictive and retrodictive science is, of course, physics. We will first address the case of classical (i.e., non-quantum) physics, and then discuss the complications that arise when quantum phenomena are introduced.

The dynamical behaviour of a classical physical system is assumed to be determined by a theory: a set of mathematical relations (typically, differential equations) describing how the *state* of the system will vary with location in space and time; and we may then use this theory to track the evolution of the system from one moment of time to the next.

More specifically: if we have sufficient information about the state of the system at some time t_N, we may input the information into the theory, and use this to make a *prediction* of the system state at any later time t_F. Moreover, in some (but not all) cases it is even possible to reverse this procedure of prediction: that is to say, when sufficient information about the state of the system at the time t_N is available, we may similarly use the theory to infer the state at the earlier time t_P. This is what we have referred to as *retrodiction.*

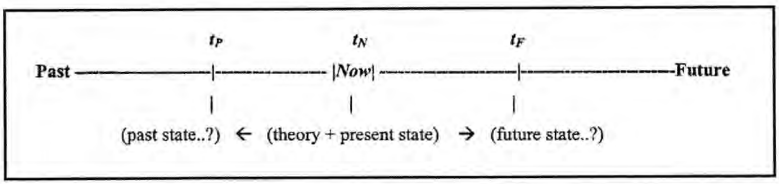

To illustrate, let us consider briefly two concrete examples, both as treated by the dynamical theory of Newtonian mechanics:

1. As noted in the introduction, we can observe the present position and velocity of the Earth and the Moon in the Solar system, and then input these data in the theory to determine with great accuracy when the next eclipse of the Sun will occur, and where (i.e., in which

regions on the Earth's surface) it will be visible. Moreover, it is possible to use the same observed initial data and dynamical theory to determine similarly when and where the previous eclipse was observed. Hence, in this particular application of Newtonian mechanics both prediction and retrodiction are possible.

2. (A more down-to-earth example, considered in a little more detail.) A material body with initial position x_N on a *horizontal plane surface* is, at a time t_N, set sliding in rectilinear non-rotating motion in the plane, with some given initial velocity. Due to surface friction it will then move a certain distance along a straight line, gradually slow down and eventually come to rest in some point x_F at a later time t_F. (We assume that the only horizontally acting force on the body is that of friction between it and the plane surface.) Here, we can determine by observation various parameter values at the start of the motion (body mass, frictional coefficient, initial position, and velocity) and input them into the theory to determine where and when the body will come to rest. However, we can *not* turn this the other way around: that is, use observations made at the point x_N after the body has stopped moving, to determine the starting point x_P and initial time t_P of the motion. Thus, in this case prediction is possible, but retrodiction is not!

Why does the first example (motion of bodies in the Solar system) allow for retrodiction, while the second example (a body sliding on a horizontal plane) does not? One way to explain this: In the first case, mechanical (i.e., kinetic + potential) energy is conserved; there is *no dissipation*. In the second case, there are dissipative (frictional) forces present, which convert mechanical (here: kinetic) energy into heat. And when there is dissipation present, entropy increases and *information is lost*: observations performed at the point x_F at some later time $t > t_F$ (i.e., after the body has come to rest) cannot give us sufficient data to enable us to deduce the start position x_N and the initial time t_N of the motion. Put another way: it is not possible to observe data at the time t that enable us to deduce how long the body has been lying still in x_F, or from what direction it originally came.

We conclude that retrodiction, in the strict sense – that is, a *precise determination* of the past, using observations made in the present – is in general not possible in physics, at least not for systems and processes that involve dissipation. Note, however, that this does *not* mean that we cannot track the development of such systems backwards in time. We can in fact do that, but we then need to introduce *additional assumptions* about the past history of the system. In the example

discussed above, we may make such assumptions: how long the body has been lying at rest in x_F, from which direction it arrived at x_F, and so forth – and then use them, together with the theory, to obtain an estimate of where and when the motion started. And the credibility of this estimate is then contingent on the plausibility of the additional assumptions.

This last example is, of course, rather simplistic and contrived; it is chosen such as to illustrate the *theoretical* possibilities of prediction and retrodiction. In actual practice, however, the situation is often a lot more complex: For one thing, a determination of the initial state is never exact: there is always some inaccuracy attached to the measurement of initial parameter values; and this, of course, leads to uncertainty in the predicted values that characterise the final state of the system.

Indeed, some of the initial parameters may be very difficult (or even impossible) to measure directly; we may have to estimate them by indirect arguments. Moreover, the theory may be very complicated to handle, and require a large amount of computing power.

However, in the examples discussed here we are not concerned with such practical considerations – our goal is to clarify under what circumstances prediction and retrodiction are possible *in principle*.

Now, let us examine some more realistic situations: First, we turn to the science of geophysics, and consider the hypothesis of *continental drift*: This hypothesis, which now seems to be accepted by the majority of geophysicists, says that the large land masses of the earth (the continents) are floating on a base of fluid magma, which allows them to move about very slowly on the earth's surface. Indeed, it is hypothesised that many (perhaps all) of the continents were long ago joined together in one large super-continent, but that they have at some earlier stage split apart from each other, and in the course of time drifted to their presently observed positions. (As evidence of this past joining of the continents, one may for instance point to the remarkably good fit that is exhibited between the coastlines of eastern South America and western Africa.)

The description of such large-scale motions belongs to the theory of hydrodynamics, and one may ask: based on this theory and on present observations, what can we infer about the positions of the continents, in the future and in the past? (Among such observations are different kinds of seismic activity, generally taken to arise from tectonic plates in relative motion.) As it turns out, a lot more interest and effort has gone into hypothesising where the continents were in the past, than into where they will go in the future.

However, various relevant physical parameters, such as fluid properties of the magma on which the continents are floating, are difficult to measure in practice; their values have to be inferred from indirect and rather uncertain arguments, and this leads to corresponding uncertainty in the estimation of both past and future states.

Moreover, continental drift is a highly dissipative process; thus, there is not even a theoretical possibility of retrodicting in a rigorous way the past distribution of continents from the present. Still, various additional assumptions have been proposed, and their merits and plausibility argued between geophysicists; and these assumptions have then, together with the theory and present-day observations, led to various models of what the surface of the earth did look like, in different epochs of the past. One might perhaps call this an example of *weak retrodiction.*

Second, we can look at the difficulty of making accurate *weather predictions,* as alluded to in the beginning of this paper. (One might also discuss the possibility of weather *retrodiction*; however, this would seem to be a rather pointless exercise, and will not be considered here.) The problem with such prediction is that theories describing weather systems are *chaotic.* This means, among other things, that final conditions depend very sensitively on initial conditions: quite small changes in the observed quantities that define the initial state can lead to dramatic changes in the end state. Hence the *accuracy* of the initial conditions is crucial to the prediction. In principle we would need to measure the initial parameters (such as, for example, the present temperature and air pressure of the atmosphere) *exactly,* and *at every point,* in order to make an exact prediction of the weather at some later time. In practice, this is of course not possible; we have to choose some level of accuracy of measurement, and measure only at selected sample points. But then the resulting weather, as predicted to occur at a later time, depends critically on this choice of the measurement accuracy and the grid of sample points. To conclude: the dynamical theory does indeed allow us to predict the weather – but it is not possible to determine the *present weather* with sufficient accuracy to ensure that the *predicted weather* is the weather that will actually happen!

This, of course, does not prelude attempts to estimate future weather – in fact, such estimates are offered worldwide, and every day. However, these estimates are then not to be classified as *prediction,* in the exact sense defined here, but rather as *forecasting*: that is, as qualified guessing, making use of the initial data at hand, and of increasingly

sophisticated mathematical techniques (as well as massive amounts of computing power) in the processing of these data. Indeed, the quality of such estimates is improving, but (as we all know) the forecast of tomorrow's weather can still give us some surprises. This might be referred to as *weak prediction.*

As a third example, we consider *cosmology.* This discipline is generally regarded as a branch of physics. It attempts to describe the Universe as a whole, or at least to explore the meaning of such a notion, using physical theory. One major goal of cosmology is to extrapolate presently observed conditions backward and forward in time as far as possible, asking questions such as:

- Did the Universe begin its existence at some definite time in the past (in the standard model of cosmology often referred to as the Big Bang), or has it existed forever?
- And will it end its existence at some definite time in the future (similarly designated as the Big Crunch), or will it go on existing forever?

This is, of course, prediction and retrodiction on a truly grand scale – perspectives that indeed fire the imagination of us all, whether or not we are scientifically trained. Clearly, they involve the *long-range* time windows as viewed from the moment *Now* (see the figure): we are here considering prediction of the far future, that certainly cannot be verified, and retrodiction of the remote past, well beyond any possibility of human experience. To do this, it is necessary to observe initial conditions here and *Now*, and use the dynamical theory to extrapolate these conditions into the remote past or remote future. A great number of books and articles that address these issues, ranging from popular to quite advanced, may be found in the literature.

To conclude, classical physics does in principle allow theoretical prediction of future conditions, and (sometimes) also retrodiction of past conditions.

Now, let us see how this appears in *quantum physics.* Here the dynamical behaviour of a system is described by a Schrödinger equation, which determines how the *probability distribution* of all possible states of the system will evolve in space and time. This evolution of probabilities may be predicted: that is to say, given sufficient information about the state at a time t_N, we may determine the probability of any system state at a later time t_F.

However, the *actual state* of the system at time t_F is not determined by the Schrödinger equation, but by a phenomenon often referred to as a *folding (or collapse) of the wave function.* When the system is observed

at the time t_F, it will fold – that is, one of the possible system states at this instant of time will be realised, and all the others will disappear. Moreover, the *selection* of this actual state, when the system folds, is probabilistic: in a given experimental run, any of the possible states may be the one that is realised; but in a large run of repeated experiments, the frequency of actually realised states at the time t_F will be proportional to their probability at this time, as determined by the Schrödinger equation.

As a very simple example of such a system, we may consider a given material sample of some radioactive element, with a half-life T. This means that after a time period T, half the atoms in the sample will have decayed. However, if we look at any individual atom, we cannot predict when this atom will decay – it could happen in the next ten seconds, or sometime tomorrow, or after a thousand years! But if we consider a very large number of atoms (such as there will be in any macroscopic material sample), the *average decay time* will be T – that is to say, at the end of this time period half of the atoms will have decayed, and the other half will not.

This very brief and condensed exposition of quantum physics may serve to illuminate a fundamental conceptual issue of this theory, which has led to endless discussions among physicists and philosophers of science. In simple terms, the problem is this: How does the system know, when observed in any one particular experimental run, which concrete state it should fold into? To see why this is a problem, we may note that the concept of probability that appears here is very different from that which appears in classical physics – even though it is handled by a similar mathematical formalism. Let us consider the issue in some more detail.

The notion of "classical probability" is, in essence, a *makeshift strategy*: probabilistic and/or statistical methods are used in situations where we do not have sufficient detailed knowledge of the system to be able to perform exact calculations. However, it is implicitly assumed that such knowledge may always (in principle, if not in practice) be obtained, with the consequent possibility of making exact predictions.

Consider, for instance, a *tossed coin*. If we were able to determine, with sufficient accuracy, the necessary physical parameters at the moment of the toss – the initial velocity and spin rate of the coin as it leaves the hand, the air resistance, the vertical distance from the ground, the elastic properties of the ground surface (which may allow the coin to bounce when it hits the ground), and so forth – we could in principle *predict*, using the theory of Newtonian mechanics, its final

state (here, whether it will end up heads or tails). In the real world, this is of course totally impracticable, due to both the difficulty of making sufficiently accurate initial measurements and the extreme complexity of the necessary computations.

But it is *possible in principle*: the initial parameters are assumed to possess definite values (although measuring them is not a realistic proposition), and an input of these values into the dynamical theory would then allow for prediction (though the computational work would be prohibitive). Hence the motion of the coin is determined: that is, its final state (heads or tails) is determined by its initial state. The random character of the toss is thus a consequence of our ignorance of the initial parameters: if we had known their values, we could have predicted the outcome of the toss, and thus have no necessity to resort to probabilistic considerations.

Similar arguments can be made when studying, say, the dynamical behaviour of a given *macroscopic volume of gas*, within the framework of classical physics. The number of gas molecules is very large, so that it is certainly not practicable to keep track of the motion of each individual molecule; hence, we have to fall back on the theory of statistical mechanics, and define various macroscopic quantities (such as the pressure and temperature of the gas) in terms of molecular properties (such as masses and average velocities), using Newtonian mechanics. The velocities of individual gas molecules are never measured – but it is assumed that these velocities are there, and could in principle be determined by measurement, if we should wish to make the effort to do so. Thus the same applies here as for the coin. If we were to gather sufficient information of the initial state of the system (which is theoretically, though certainly not practically, feasible), and use the theory to perform the necessary computations (which would in any case be prohibitively complicated), we could *in principle* predict in detail its later states.

In the case of quantum theory, however, the *conception of probability* is fundamentally different. Here the use of probabilistic / statistical methods is *not* a makeshift strategy that we have to resort to, because we lack detailed information about the system. On the contrary, the probabilities that enter into the Schrödinger equation represent the *only* valid information that the theory can offer, concerning the dynamical development of the system. In other words, there do not exist any hidden variables that we could (in principle, at least) measure, and that would enable us to predict the state that is realised when the system folds. Thus, with reference to the example of radioactive decay that was

discussed above, there are no physical properties of the material sample that we can observe, which would help us determine when any one particular atom is going to decay. This proposition is supported by the experiments of Aspect and others, which – in conjunction with Bell's inequalities – appear to preclude the possibility of any kind of hidden variables underlying the statistical treatment of quantum systems.

(In the teaching of physics, this is sometimes expressed in a somewhat anthropocentric way, by saying that the atom "decides to decay," when it "wants to do so." This is perhaps an unfortunate metaphor, since it seems to suggest that the atom possesses some volition of its own – however, it does illustrate rather well the "wilfulness" and "unpredictability" of such physical systems.)

To conclude, in quantum physics there is *no prediction* of the actual (i.e., physically realised) final state of a system, deducible from knowledge of the initial state of the system. There is a very limited kind of prediction, governed by the Schrödinger equation, of how the distribution of probabilities will develop from a given initial state right up to the next folding – and that is as far as prediction can go. Moreover, there is *no retrodiction*: once the system has folded into a particular state, no observable information about previous (non-realised) probabilities is available. Referring once again to the radioactive decay example: before the atom decays, it is not possible to predict when this will actually happen in the future; and after it has decayed, it is not possible to retrodict when this did happen in the past.

(It may be remarked that it was precisely this fundamental unpredictability of physical processes at the quantum level that Albert Einstein refused to accept, and which led him in the later years of his life to distance himself from the mainstream community of scientists, in a largely solitary – and, as it turned out, unsuccessful – attempt to establish a fully predictive science of physics.)

Finally, we ask: in those cases where physics does offer the possibility of prediction, what *epistemic* status can be attributed to this "knowledge of the future?" The position of radical constructivism on this issue is that such knowledge, like any other, is *constructed* – here: on the basis of the theory, together with present observations and (possibly) additional assumptions. When this constructed knowledge refers to events in the *short-time* future window, it can be verified by observation, as noted above; and such verifications may then support the theory – that is, strengthen its status as viable knowledge in the minds of the knowers. For events in the *long-time* future window, such verification is not possible; here we can only say that this knowledge has been

constructed in such a way as to be compatible with present-day theory and experience.

Retrodiction – The Historical Perspective

The physical sciences (physics, astronomy, cosmology) are special, in that they are so strongly based on a formulation in mathematical terms. One consequence of this is, as we have seen, that the theory will allow (in some cases) a *rigorous* prediction of future events, or retrodiction of past events. The situation is rather different for the biological and geological sciences: here the theories will generally be of a more descriptive and qualitative character, and rigorous arguments will be correspondingly rare. Both predictions and retrodictions then tend to be more tentative: given present observations, the theory may say that such-and-such events have some *likelihood* of happening sometime in the future, or of having occurred sometime in the past.

Many research areas (belonging to natural and other sciences) are predominantly *retrospective* – that is, they are mainly involved with a study of the past, and little (if at all) concerned with future events. The goal of such retrospective sciences is to establish plausible hypotheses about events as they may have happened in the past, using data that we can observe in the present. The basic general requirement is then that these hypothetic past events must be compatible with present data: that is, it must be possible (using the theory) to suggest credible processes, whereby the hypothesised past conditions can *evolve in time* into the conditions that are observed now.

Past events, as studied by the various retrospective sciences, may be naturally grouped into domains, spanning different past time windows. For each such domain, there are certain sources of information (as observed in the present) that are relevant for the kind of past events that are being studied. Some examples:

- The *social sciences* (psychology, sociology, political science, jurisprudence) are largely concerned with events that have occurred "within living memory" – roughly speaking, events that may (in principle) have been experienced by persons living today. Information sources are one's own personal memories, oral communications, written records, and so on – all of events located in the short-range past time window.
- *History* is engaged in the chronicling of human society. Information sources: mainly written documentation and surviving artefacts. The time window extends back some millennia.

- *Evolutionary biology* deals with living organisms and their development in time. The main information source is fossil evidence of various kinds. Time window: some millions of years.
- The *geological sciences* (geology, geophysics) study the behaviour of large land masses. Information sources: present-day distribution, seismic activity and mineral composition of the land masses. Time window: some billions of years.
- *Cosmology* studies the past behaviour of the whole Universe (or large portions of it). Information sources are mainly observations of incoming radiation: spectra of stars and galaxies (indicating positions, velocities, and material composition), and various types of cosmic radiation. The extension of the past time window: whatever the theory may hypothesise for the age of the Universe.

This list makes no claim to completeness; and (needless to say) the past domain boundaries are not sharply defined. Within each domain, the "information density" tends to diminish with age: that is to say, the longer ago the events, the smaller the amount of relevant observational evidence today.

One might say that the purpose of retrodiction is to determine past events, in such a way as to be compatible with the events that are observed in the present. This is analogous to the purpose of prediction (determining future events to be compatible with presently observable events). But there is one essential difference. The future is, in the perspective of any given science, explicitly taken to be *indeterminate* – that is to say, we can exercise our free will to influence the course of future events, as was discussed above. The past, on the contrary, is generally assumed to be *fixed*, in the sense that its course of events cannot be influenced by present-day actions. In other words, the assumption is that in the past events have already happened in some way, and this past course of events cannot be changed by actions taken here and now.

This raises the question of whether it is possible to obtain *true knowledge of the past* – in other words, to establish what "really happened" at some time prior to the present.

For *the short-range* time window, at least, this question might seem at first glance to be trivial: surely, the faculty of memory (suitably aided by written records, tapes, observational verification, etc.) can provide each of us with such information? Indeed, this assumption is basic to any *legal proceedings*, such as in a court of law: the true facts of the matter at hand are there to be found, and the goal of the court is then to

ascertain these facts (insofar as this is practically possible), in order to arrive at a fair ruling.

However, this notion of "the true facts" is a highly problematic one, even in jurisprudence – and not only for the obvious reasons: memory may fail, or given testimony may be mendacious. In fact, our recollections of past events, as well as sensory perceptions of present events, are continually being processed in the mind by the faculty of *reflection*: a complex and many-faceted activity.

For instance, even with the best intentions of honesty, there will be some unavoidable (subconscious) *filtering* going on: how do we want to remember it, what items would we rather forget about? Moreover, there will be a *selection*: since one cannot recall (and report on) everything that has happened, it is necessary for the mind to pick the important or relevant features – and disregard whatever is considered to be unimportant or irrelevant. And finally, any kind of input to the mind, whether of present sensation or past recollection, needs to be *interpreted*: what meaning do we read into written or spoken testimony, and how do we construe – directly (as perceived) or indirectly (as remembered) – the intentions, attitudes, and actions of other agents?

In this way – using perception and reflection, both as performed before the moment *Now* – we construct our knowledge of what happened before this moment. In a word: *we construct our own past.*

This is a basic premise of radical constructivism; and it implies that the question of "what really happened" in the past is, from the point of view of cognition, epistemically meaningless! All we can do is construct our cognitive knowledge of the past to be as compatible as possible with present-day experience, and with whatever theory and other assumptions we may subscribe to – in other words, construct it to be *viable* knowledge (Glasersfeld, 1995).

Summing up: The theory of radical constructivism advocates a *relativist* epistemic position, on the status of our knowledge of both the future and the past. Such knowledge is (and must be) *constructed* by the knowers, in the context of:

- present-day perceptions and reflections,
- available recorded information, and
- choice of theoretical framework,

in such a way as to be viable in this context.

In particular, it should be emphasised that, within the framework of radical constructivism, there is no conception of obtaining true knowledge of the past; and hence it becomes meaningless to look for an account of what really happened back then.

REFERENCES

Glasersfeld, E. von. (1981). The concepts of adaptation and viability in a radical constructivist theory of knowledge. In I.E. Siegel, D.M. Brodzinsky, & R.M. Golinkoff (Eds.), *New directions in Piagetian theory and practice.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Glasersfeld, E. von. (1984). An introduction to radical constructivism. In P.Watzlawick (Ed.), *The invented reality* (pp. 129). New York, NY: Norton.

Glasersfeld, E. von. (1989). Constructivism in education. In T. Husén & T.N. Postlethwaite (Eds.), *The international encyclopaedia of education, Supply.* (Vol.1, pp. 162-163). London, UK: Plenum Press.

Glasersfeld, E. von. (1992). A constructivist view of learning and teaching. In R. Duit, F. Goldberg, & H. Niedderer (Eds.), *Research in physics learning. Theoretical issues and empirical studies* (pp. 29-39). Bremen.

Glasersfeld, E. von. (1995). *Radical constructivism: A way of knowing and learning.* London, UK: The Falmer Press.

Glasersfeld, E. von. (2000). Problems of constructivism. In L.P. Steffe & P.W. Thompson (Eds.), *Radical constructivism in action: Building on the pioneering work of Ernst von Glasersfeld* (pp. 3-9). London, UK: Routledge/Falmer.

Andreas Quale attained his Ph.D. at the University of Oslo, Norway in mathematical physics. He has taught upper secondary school, college, and university courses in mathematics, physics, and computer science (programming). Recently he has been responsible for the Computer Science Teacher Training Program at the University of Oslo. He has authored (in Norwegian) textbooks in physics and computer science for upper secondary school and college level students and has published numerous papers on science teacher education and philosophy of education.

Author's Address:
Dept. of Teacher Education and
School Development (ILS)
University of Oslo
NORWAY
EMAIL: andreas.quale@ils.uio.no